NPS52-86-020

NAVAL POSTGRADUATE SCHOOL
Monterey, California

THESIS

DESIGN GUIDELINES FOR A RULE-BASED
PASSIVE SURVEILLANCE SYSTEM

by

Kirk Edward Jennings

September 1986

Thesis-Co Advisor:	Neil C. Rowe
Thesis-Co Advisor:	Calvin R. Dunlap

Approved for Public Release; Distribution Unlimited

Prepared for: Commander, Space and Naval Warfare Systems Command
(PMW 180 - 421)

NAVAL POSTGRADUATE SCHOOL
Monterey, California

Rear Admiral R. C. Austin
Superintendent

David A. Schrady
Provost

The work reported herein was supported in part by the Defense Advanced Research Projects Agency under ARPA Order No. 5141, N6227186WR60080.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1a REPORT SECURITY CLASSIFICATION UNCLASSIFIED	1b. RESTRICTIVE MARKINGS
2a SECURITY CLASSIFICATION AUTHORITY 2b DECLASSIFICATION / DOWNGRADING SCHEDULE	3 DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; distribution unlimited.
4 PERFORMING ORGANIZATION REPORT NUMBER(S) NPS 52-86-020	5 MONITORING ORGANIZATION REPORT NUMBER(S)
6a. NAME OF PERFORMING ORGANIZATION Naval Postgraduate School 6b OFFICE SYMBOL (If applicable) 52	7a NAME OF MONITORING ORGANIZATION Naval Postgraduate School
6c. ADDRESS (City, State, and ZIP Code) Monterey, CA 93943	7b. ADDRESS (City, State, and ZIP Code) Monterey, CA 93943
8a. NAME OF FUNDING / SPONSORING ORGANIZATION Commander, Space & Naval Warfare Systems Command 8b. OFFICE SYMBOL (If applicable) PMW 180-421	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER
8c. ADDRESS (City, State, and ZIP Code) Washington, DC 20362	10 SOURCE OF FUNDING NUMBERS PROGRAM ELEMENT NO: DARPA N62 PROJECT NO: 27186WR60080 TASK NO: WORK UNIT ACCESSION NO:

11 TITLE (Include Security Classification)

Design Guidelines for a Rule-Based Passive Surveillance System

12 PERSONAL AUTHOR(S)

Kirk Edward Jennings in conjunction with Neil C. Rowe and Calvin R. Dunlap

13a TYPE OF REPORT	13b TIME COVERED	14 DATE OF REPORT (Year, Month, Day)	15 PAGE COUNT
Masters Thesis	FROM Sept '85 TO Present	September 1986	52

16 SUPPLEMENTARY NOTATION

17 COSATI CODES			18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	Artificial intelligence; Rule-based system; Passive Surveillance System; blackboard architecture; frames data structure, uncertainty

19 ABSTRACT (Continue on reverse if necessary and identify by block number)

This paper addresses the application of artificial intelligence to passive surveillance systems that use waveform analysis as their primary means of detecting, classifying and locating a specific target. Discussion is further limited to those passive surveillance systems which must deal with considerable noise in the data.

Present methods, which use visual examination of the waveform data for the detection of target waveforms, is complicated, time consuming, and requires considerable expertise. The lack of prior knowledge of the nature of the noise, (e.g., frequency spectra, amplitude, or dynamics), means that the majority of signal analysis must be done by experts.

This study discusses and recommends a rule-based system which uses the following artificial intelligence structures: the blackboard architecture, and the frames data structure. Sources of uncertainty are also discussed and methods of dealing with it are suggested. This study recommends that the symbolic representation language be carefully

20 DISTRIBUTION / AVAILABILITY OF ABSTRACT ☒ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT ☐ DTIC USERS	21 ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED
22a NAME OF RESPONSIBLE INDIVIDUAL Neil C. Rowe	22b TELEPHONE (Include Area Code) (408) 646-2462 22c OFFICE SYMBOL Code 52Rp

DD FORM 1473, 84 MAR — 83 APR edition may be used until exhausted. All other editions are obsolete

SECURITY CLASSIFICATION OF THIS PAGE

selected for conciseness, efficiency, and a vocabulary rich enough to express everything desired by the experts. A learning knowledge source is also recommended.

Design Guidelines
For a Rule-Based
Passive Surveillance System

by

Kirk E. Jennings
Civilian, Naval Ocean Systems Center
B.S., University of Hawaii, 1968
M.S., University of Hawaii, 1969

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN COMPUTER SCIENCE

from the

NAVAL POSTGRADUATE SCHOOL
September 1986

ABSTRACT

This paper addresses the application of artificial intelligence to passive surveillance systems that use waveform analysis as their primary means of detecting, classifying and locating a specific target. Discussion is further limited to those passive surveillance systems which must deal with considerable noise in the data.

Present methods, which use visual examination of the waveform data for the detection of target waveforms, is complicated, time consuming, and requires considerable expertise. The lack of prior knowledge of the nature of the noise, (e.g., frequency spectra, amplitude, or dynamics), means that the majority of *signal analysis* must be done by experts.

This study discusses and recommends a rule-based system which uses the following artificial intelligence structures: the blackboard architecture, and the frames data structure. Sources of uncertainty are also discussed and methods of dealing with it are suggested. This study recommends that the symbolic representation language be carefully selected for conciseness, efficiency, and a vocabulary rich enough to express everything desired by the experts. A learning knowledge source is also recommended.

TABLE OF CONTENTS

I. INTRODUCTION

A. PROBLEM DEFINITION AND STATEMENT

1. Delineating the Problem

This paper addresses the application of artificial intelligence to passive surveillance systems that use waveform analysis as their primary means of detecting, classifying and locating a specific target. Discussion is further limited to those passive surveillance systems which must deal with considerable noise in the data.

It is not the intent of this paper to present methods of detection, classification or tracking, but to outline artificial intelligence structures that could be used in a computer to support the solution of some waveform analysis problems.

2. Origin of the Material Presented

The thoughts and ideas presented here were gained by studying material written on the subject, personal experience in related fields, and seven weeks of direct observation of some very talented, and dedicated surveillance experts who work very hard to get any results at all. It is sincerely hoped that the guidelines presented in this paper will aid in the design of a system that will be of some benefit to them.

3. Background and Introduction to Waveform Analysis

a. Human Expert Verses Artificial Intelligence

It is the opinion of the author that the mere application of artificial intelligence to the problem will not solve the problem. The

human expert, intelligently employing the artificial intelligence system as a tool, will solve the problem.

b. Limitations of the Examples Used

Examples presented in this paper are suggestive of the type of tasks, relative to the problem, that the artificial intelligence system can do.

c. Definition of Waveform Analysis

Waveform analysis in general means studying the frequency components of that waveform. This is usually done by transforming the waveform, which is a composite of many frequencies, into its spectral components and studying them. Signal processing devices are employed to transform the waveform into a format which can be analyzed by either man or machine.

One has only to study the theory of radio and sonar to understand why only signal processing devices are needed to do the radio waveform detection and a *man*, primarily, is required to do the waveform detection in sonar.

Algorithms exist which give satisfactory results in the theory of radio waveform detection. The parameters involved are known. In sonar systems, the principles involved are understood but it is virtually impossible to account for and measure all the variables. Signal processing devices can do the part of the job where algorithms exist but the rest falls to man.

d. Distinction Between Detection and Classification

A distinction is made between detection and classification. The former implies that an object of the same type or class, as the target, is present in the data. Classification implies identification of an instance of that class, i.e., the object.

Tracking is accomplished by locating the targeted source and tracing its movement as a function of time.

4. Unsolved Problems in Waveform Detection

a. Waveform Problems With No Satisfactory Solutions

A difficult problem in the field of waveform detection, which remains of interest, is that of detecting target signals in situations where the amplitude of the signal is low with respect to the noise. This occurs where, for one reason or another, the source signal level cannot be increased relative to the noise, and the targeted signals can only be monitored passively.

(1) Problem Has No Complete Algorithms. This type of problem cannot, in general, be completely described by an algorithm for one or more of the following reasons:

* Sensors cannot be put close enough to the target source, thus, low signal amplitude at the sensors

* Noise sources are as loud as the target

* Noise sources exist which are similar in frequency spectra

* Time variant environmental conditions, for instance; the transmission media could vary in density around the sensor causing the waveform to travel faster or slower

(2) Surveillance Example. An example of this type of problem could be trying to pick out the sound of one particular animal and track it in a jungle that may be filled with the sound of other animals.

The sensors, located in various parts of the jungle, are extremely sensitive and can monitor the sounds that the animals make in most locations. The trees and rocks sometimes alter the sound that comes to the sensors. Noise from other sources is louder, at times, than the animal itself. The animals are free to move around and sometimes go out of sensor range.

Signal processing devices alone cannot do the complete job of detection. A human expert, correlating his knowledge of the animals vocal emissions with that of the waveforms from the field sensors, and with his innate understanding of animal behavior, will have the greatest chance of successfully detecting and tracking an animal.

5. Artificial Intelligence Suggested to Aid Expert

This paper suggests the use of artificial intelligence in the form of a rule-based system to aid the human expert in studying these types of problems. Recommendations are made to design a rule-based system which uses the blackboard architecture as a control structure, and the frames structure for pattern matching and knowledge representation. A learning knowledge source is recommended for automatic knowledge acquisition.

6. Unique Combination of Artificial Intelligence Structures

Eight well known rule-based systems in use today are discussed by Hayes-Roth, Waterman, and Lenat in (Hayes-Roth, 1983, pp. 169-215). The combination of artificial intelligence structures proposed by this paper has not, to the author's knowledge, been brought together and used on the type of passive surveillance problem being discussed.

7. Topics Covered in the Paper

Section II discusses present methods used by the experts for detection, and ways an artificial intelligence computer system could be applied. Several artificial intelligence techniques, which support human experts investigating problems of this nature are presented in Section III.

II. DETECTION OF WAVEFORMS IN NOISE USING PASSIVE SENSORS

A. APPLICATIONS

Several areas in which passive surveillance systems could be used are:

* Monitoring seismic activity using arrays of seismic transducers.

* Avalanche detection done by using arrays of audio transducers.

* Detection, and tracking of objects that emit sound energy such as: vehicles, animals, etc., using arrays of appropriate transducers.

* Detection of radio sources in the sky using radio telescopes.

B. PRESENT METHODS

1. Visual Presentation of the Data

Examination of the data, in the class of problems discussed in this paper, is generally done visually. The human expert is easily capable of recognizing patterns, gestalt, in situations where the target signal strength is clearly visible above the noise.

In situations where the expert has prior knowledge of the spectral waveform of the target, discussed in the next paragraph "Pattern Matching", the incoming waveform can be processed and presented in a spectrographic format. To accomplish this, signal processing devices which emulate the Fourier transform are employed to break the waveform into its component frequencies.

a. Types of Displays

This data can then be presented visually, in a waterfall like display, or on a cathode ray tube in a display called an *A scan*.

(1) The *A Scan* Display. The *A scan* display presents frequency information on the horizontal axis, and the amplitude of those frequencies on the vertical. One drawback of the *A scan* display, in this type of problem, is that the observer cannot easily see the waveform's past history. It is the past history which allows the observer to uses his eyes to integrate and filter the signal. This technique has proven to be a tremendous aid in detection.

(2) The Waterfall Display. The waterfall display solves this problem. The waterfall display presents spectral information as a series of dots on a stripchart recorder with the amplitude of a frequency component reduced to the gray shade of a dot. Each time new data arrives a new line of dots are recorded on the stripchart paper. This type of presentation of the data allows the observer to not only examine the current data but past data as well.

2. Pattern Matching

The human expert is good at picking out patterns, even in a background of noise where conventional signal processing devices alone would have difficulty. A human observer can often detect a signal despite the presence of a large amplitude noise spike, that is at or near the frequency of interest. This is done, perhaps, by ignoring the noise, or looking at the signal before and after the disturbance, or some other strategy.

a. Prototyping of the Targeted Signal

The most common method used by the expert to find target signals is to examine the display and compare it to past detections. So, prior to starting the system, recordings of the target waveform are made

in a high signal to noise environment. This *clean* signal is then displayed in many formats, studied, numerically analyzed, etc., by experts.

b. Distinctive Aspects Chosen for Prototype

This recording is analyzed and the characteristic frequency components are included in the prototype of the target. This prototype, sometimes referred to as a template, is subsequently used as a basis for comparison with the waveforms taken in the field.

c. Example of Making a Template

In the jungle example, the goal is to track a particular animal. The expert would make recordings of the animal by itself, with no noise sources. Similar recordings of other individual animals could be made and used in cases where discrimination was difficult. The sound structure from these recordings could then be carefully studied, and its peculiarities noted. The knowledge would then be used as a template to match future occurrences of the sound where the signal to noise ratio may not be ideal.

d. Expert Uses Other Attributes of the Target

In hard cases where noise is dominant, the expert tries to bring in other characteristic parameters that the targeted signal may possess to substantiate or disprove his beliefs. An example of this would be, if the source of the signal were dynamic, that is it moves, then the expert may use his knowledge of the limitations in the targets mobility to eliminate false targets. Another characteristic which might be exploited is a knowledge of the topography, walls, buildings, mountains, other obstructions. Knowledge of this kind could be used as evidence that the target source could or could not be emanating from a specific location.

e. Much of the Experts Knowledge is Stored in His Head

The expert uses different strategies based on visual observations and intelligence inputs. Intelligence inputs may be in the form of information regarding weather conditions; or that all of the animals have migrated out of the jungle; etc.. Much of this type of knowledge and problem solving strategy which the expert uses is not stored on paper, and is not strictly related to waveform analysis, but is the sum of all his life experience, education, and training. An artificial intelligence system with a good natural language interface can be used to capture this expertise.

f. Some Cases Do Not Require an Expert

Depending on the appearance of the signal, there may be no need to use other strategies. The strength of the signal may be so high, with respect to the noise, that detection is easy.

g. Mechanical Templates Used in Pattern Matching

In some cases the patterns found on the display may be so recurrent and easily classified that physical devices, also called templates, are made and used to compare with the displayed data.

3. Use of Signal Processing Devices

Waveform analysis is generally considered a signal processing task, done with numerical routines. This is especially true in cases where there is a high signal to noise ratio. Signal processing, in cases like this, is generally accomplished in real time using devices such as - narrow band filters, and phase locked loops. These devices can be programmed or built to track frequencies within a given bandwidth. The phase locked loop is

capable of tracking a single frequency even when that frequency is drifting.

a. Signal Processing Devices Used to Alert the Expert

The expert can, to a limited extent, use these devices as an alerting signal to indicate the presence of one or more signals. These devices must be selected by the expert for each task. If the circumstances of the problem change, these alerting devices must be redirected. This may mean devices that are pre-set by the factory, or ones that can be set in the field. The artificial intelligence computer can be used to change the parameters to signal processing devices if those devices are designed for that type of interface.

C. SOURCES OF UNCERTAINTY

1. Noise in the Data

Partially incorrect evidence is common in systems that rely on sensory data. Data from the sensors is determined to be noise if it is not associated with the target signals. Noise can cause uncertainty about whether the evidence is relevant or not, so before the expert can deal with the data, he must first deal with the uncertainty of that data.

a. Eliminate Noise Sources to Reduce Uncertainty in Data

One strategy the expert might use to reduce the uncertainty in the data is to reduce the clutter of the display by identifying noise sources. He could then filter them out, physically or mentally, as not being potential target signals. Potential problems can arise because the signal observed is close to the target signal or the origin of the noise cannot be determined. Problems such as these make this a difficult technique to use.

b. Localize Noise Source to Reduce Uncertainty in Data

Another strategy which might be employed, in the case where the noise is not identifiable, is to try and localize it to an area which the target could not be. This technique implies the use of directional sensors. This would mean cross-correlating the signals from the sensors or using other techniques in order to pinpoint the origin of the signal.

Cross correlation is then done by trying to match data on one sensor to similar data on other sensors. This alternate strategy may be possible, but again, it is difficult, time consuming, and generally not done because little time can be spared from tracking the primary target.

c. Example of Reducing Uncertainty by Localization

In monitoring the waveforms from the jungle there might be, at any given moment, airplane activity, atmospheric disturbances, or other unknown noise sources. The expert could scan the data, see several noise sources which he recognizes, and by some means remove them from the display. Once this is done he is looking at a less cluttered display, though some unknown noise sources may still remain. Working with the second strategy, using localization, the expert might eliminate noise sources which appear to be coming from places outside the jungle, or where the animal could not be.

D. DISADVANTAGES OF PRESENT METHODS

Several disadvantages have already been mentioned but are restated for emphasis. The present methods used for the detection of targeted waveforms using passive sensors are complicated and time consuming. The lack of prior knowledge of the nature of the noise, (e.g., frequency

spectra, amplitude, or dynamics), means that the majority of the signal analysis must be done by the experts.

The expert may have other commitments. It may be difficult to find an expert, or there may only be one in the world. Perhaps the volume of signals which should be studied by the expert exceeds his capacity to process. These are only a few reasons for trying to record the experts knowledge.

1. Potential Loss of Expert Knowledge

A second-order effect occurs in the area of training new expert observers to take the place of tiring, retiring, or transferring experts. Because experts in the field cannot always explain exactly how they accomplish what they do, some of the corporate knowledge leaves, and the expertise is lost or must be re-discovered.

III. ESSENTIAL CHARACTERISTICS FOR THE RULE-BASED SYSTEM

A. NEED FOR ARTIFICIAL INTELLIGENCE

1. Signal Processing Cannot Remove All Noise From Waveform

Signal processing techniques alone cannot perform the entire task of signal detection in the presence interfering noise sources. It cannot because there is always uncertainty if *all* noise cannot be identified and removed from the waveform. Signal processing devices have no way of identifying noise sources on their own, let alone the formidable task of removing them from the waveform.

2. Humans Can Act As Adaptive Signal Processors

Human experts have the ability to adapt to a wide spectrum of circumstances and events. They can focus on signals which are of interest, and can use other parameters associated with the target and noise sources to achieve better results. A symbolic computer should be able to do some of the reasoning tasks the expert does, do them well, and do them faster.

3. Artificial Intelligence Aids Expert in Symbolic Reasoning

This type of problem is a good candidate for artificial intelligence because it often involves more than one domain expert. The experts may use one or more strategies involving: pattern matching, guesswork, heuristics, etc., particular to their expertise. The frames data structure, discussed later in Paragraph C., in "The Frames Data Structure", is well suited to store and use this type of knowledge. This data structure allows

the domain experts to write in a natural language, and get understandable feedback while editing, debugging, and tracing.

4. Help Needed in Examining Huge Amounts of Data

If data are coming in - twenty four hours a day, every day, and there are many sensors, then there is a huge amount of data to examine. Looking abstractly at the numbers, one target, or noise source can appear on all sensors, giving the appearance of multiple targets. The number of possible targets to consider then becomes the number of sensors multiplied by the sum of the target and noise sources. This number, of course, becomes large when the number of sensors increases. The problem then becomes unwieldy for the human expert by the sheer volume of data, let alone the task of separating signal from noise. A symbolic computer could reduce the expert's workload, and possibly, the requirement for many experts.

5. Strategies Used By Experts Aided By a Rule-Based System

The experts generally have studied other aspects of the target, attributes such as its maximum speed, its last independently verified sighting, how far the vehicle can go on a tank of gas, etc.. Individual experts in each field may have different ways of using these facts in solving the problem.

The expert scans the data, as discussed earlier, and may use the targets attributes to gain evidence for or against the presence of a target signal. Progress towards that end could be accelerated if many of the repetitive mental tasks, that the domain expert does relating these attributes to the target, could be done by a symbolic computer architecture that supports multiple knowledge sources.

6. Experts Sometimes Use Alternative Solutions

There are times when the expert would like to pursue several strategies because he is sure one of then will be successful. The human expert can do this but the cost in time may force him to overlook potential solutions. The symbolic computer that supports multiple knowledge sources can pursue alternative strategies rapidly. This is discussed further in Paragraph B., Sub-paragraph h. "The Advantages of Parallel Processing".

7. Experts Want to Understand and Easily Guide the Computer

Most experts in the disciplines that would like to take advantage of what a symbolic computer might do are not expert computer programmers. As individuals, they may be able to write in Pascal or Ada but writing programs is not time spent being a domain expert. Alternatively, if they let a systems programmer write in some esoteric programming language then modifications will always require the programmer. A rule-based system, Hayes-Roth(Hayes-Roth, 1985, pp. 921- 932), can provide the structure that facilitates the use of a natural language symbolic reasoning program. Experts can then easily input their own knowledge, and do their own editing, and debugging.

8. Symbolic Computer Captures the Domain Experts Knowledge

These domain experts can use a common language to develop programs for the symbolic computer by choosing the blackboard architecture, discussed later in this section. This has the added benefit of capturing the expert's knowledge which can then be studied, modified, and learned by others.

B. RULE-BASED SYSTEMS

1. Categories of Rule-Based Systems

Rule-based systems fall basically into several categories, Hayes-Roth, Waterman, and Lenat (Hayes-Roth, 1983, pp. 13-15), interpretation, prediction, diagnosis, design, planning, monitoring, debugging, repair, instruction, and control. The rule-based system being studied falls into the interpretation category which includes surveillance, speech understanding, image analysis, chemical structure elucidation, signal interpretation, and many kinds of intelligence analysis. A rule-based signal interpretation system tries to explain observed data by assigning to them symbolic meanings which describe the situation or system state which accounts for the data.

2. Symbolic Representation

a. Symbolic Representation Must be Carefully Chosen

Artificial intelligence tools can ease the task of development by the domain expert. As indicated by Bolc, (Bolc, 1984, p. 31), the artificial intelligence tools used in the system should make it easy for the non-programmer to communicate with the program without prior experience or training. That is, the language which the expert uses to communicate with the computer, and vice versa, must be easily understood, concise, efficient, and have a vocabulary rich enough to accomplish everything needed by the task.

b. Symbolic Representation Needed During Development

In the development stage it is essential for the knowledge and domain experts to be able to understand and shape the program's

behavior. If it is hard for the domain experts to communicate with the rule-based system, development will be slowed. Symbolic representation is needed that clearly and succinctly describes both intermediate and final results which are appropriate to the problem being investigated. For example, if *some* tiger was detected by the system which was not *the* exact one being sought, the message *vx* has little meaning, but *tiger, but no cigar* is easy to interpret by anyone.

c. Intelligible Feedback for Editing, Debugging, and Tracing

There should be an editor, and debugger which are easy for the domain experts to use. Trace facilities, which allow the experts to follow the progress of the knowledge sources as it is being debugged, should also be provided.

It is essential for the rule-based system to feedback information to the expert which tells - how it is doing what it is doing, and why. Feedback to the users query of *how* or *why* should explain how a decision was arrived at, or why it needs answers to certain questions. Trace facilities can also provide a form of feedback that can be used to train other experts. The frames data structure, of Fikes and Kehler (Fikes, 1985, pp. 904-920), discussed later, supports this type of feedback.

3. Knowledge Sources

a. Definition of Knowledge

A definition of knowledge is: the skills and abilities needed to be successful. Knowledge, in relation to pattern matching and analysis, refers to constraints and associations between objects and events occurring in the real world.

b. Definition of Knowledge Sources

A knowledge source is a grouping of production rules which form a module of knowledge. Production rules are sets of antecedent - consequent pairs. An example of a production rule could be the old whirlwind tour of Europe remark *If it's Tuesday, this must be Belgium.* In Prolog, an artificial intelligence language, the production rule would look like: *Belgium :- Tuesday.* The antecedent is Tuesday, and the consequent Belgium.

c. Contents of a Knowledge Source

Knowledge sources are made up of general facts, processes, relationships, heuristics, etc. associated with the subject. Knowledge Sources might be:

* Technical parameters of the targets
* Strategies
* Sensor peculiarities
* Propagation path anomalies
* Environmental factors
* Noise sources
* Dynamics of the target and noise sources

d. Knowledge Sources Are Large Grained Production Rules

Each knowledge source can be thought of as a large grained production rule which reacts to changes produced by other knowledge sources. A first cut at assembling the production rules for a knowledge source would be to list all of the knowledge that an expert uses, specific to a facet of the signal source, e.g., mobility, which is sometimes called the target's dynamics.

e. Knowledge Sources Can Be Sub-divided

Nothing is sacred in the definition of this type of module, further examination can lead to further subdivision. Subdivision of this set of rules might be done because another knowledge source could use the same function, or this group is active only if a certain event occurs, or some other need by the control system, such as excessive execution of the present module. Debugging these knowledge sources becomes successively simpler.

f. Advantage of Knowledge Sources At Run Time

Another advantage of having separate knowledge sources comes at the time the module is loaded into the computer. Knowledge sources, groups of production rules, can be pulled into the computers active memory area, activated, their results posted on the blackboard, and then removed from memory. The advantage being, the computer does not need to go through all the production rules in all the knowledge sources to get the results of one. Another advantage is that two different knowledge sources can use the same internal variables with no adverse interaction because they are not active at the same time.

g. Advantages During Development

Production rules that are modularized along the target's attributes and facets are much easier to understand, modify and debug, than keeping all the production rules together. Convoluted software is easy enough to develop without adding fuel to the fire, i.e., mixing all the rules together.

h. The Advantages of Parallel Processing

(1) Example of Parallel Processing. Knowledge sources can, if computer facilities allow, be run in parallel, thus setting up the mechanism for competing solutions or alternative solutions. Parallel processing can save time where alternative solutions were run in series.

(2) Non-determinism and Heuristics. For example, if the expert had two strategies which he used because if one did not succeed the other would. He might run the one that had the highest probability of success, and the next one after that if the first one failed. Parallel processors could carry out both strategies at the same time and stop the process when either succeeded.

Another advantage of artificial intelligence is that it can deal with this type of non-determinism, by using heuristics to pare down the potentially explosive number of cases which could result.

i. Multiple Knowledge Sources

There is little doubt that there will be more than one knowledge source used in the system. Each one of the goals - detection, tracking, and identification may have several. The artificial intelligence structure must be capable of supporting these and maintaining the integrity of each module.

(1) Knowledge Sources Which Could Be Used in the Example. A frame-based knowledge source, in the form of a taxonomy, could be used for the detection and identification of a tiger. It would contain instances of tigers on the bottom level; categories of cats: tigers, lions, leopards, etc., and rules which distinguish them. Another level might contain rules that distinguish cats from other animals in the area.

A tracking knowledge source could contain rules which correlated detections of identifications on separate sensors. Results would be posted, if localization were possible.

Another knowledge source could be for animal dynamics. This could access a frames taxonomy of animal mobility with distance being at the bottom level, speed and acceleration being on higher levels. Tracking information would be processed by another knowledge source and posted on the blackboard.

There could be a data base knowledge source made up of intelligence that had been gained about the animals. This could contain rules such as: If it's not on the ground and it's night then it's a cheetah. It might also contain information about independently verified sightings, date and time included, which could be used to discount the belief that the animal was in another area. A historical file with a data base of paths or trails commonly used by the animal could be kept and similarly used.

4. User Friendly

Another essential feature that goes with ease of programmability and understandability is a user friendly, interactive display. This has little to do with artificial intelligence but is an important part of making a system easy for a domain expert to use. The use of windows, drop down menus, computer type mice, etc., should receive strong consideration over the presently popular multi-keystroke approaches.

5. Symbol Manipulation With Signal Processing Devices

A prime motive for developing knowledge-based signal processing systems is, as indicated by Kopec, (Kopec, 1982, pp. 1-6), the anticipated

advantage from combining the symbol manipulation capabilities of artificial intelligence and knowledge representation, with numerical and mathematical tools of signal processing.

a. Rule-Based System Controls Signal Processing Devices

Signal processing devices, mentioned in the previous section, can be employed by the rule-based system. For example, the raw data could be digitized with an analog to digital converter, with consecutive samples stored over a fixed period of time. The stored data could then be run through a spectrum analyzer. One or more phase locked loops could be employed, at strategic frequencies, to alert the rule-based system that a targeted signal was there. This signal from the device is sometimes called an *event* in artificial intelligence. The rule-based system could then, depending on how it was programmed, check to see if there was an *event* like this in the previous cycle. This information could then be used to build evidence that there was a spectral line which occurred over a length of time. The expert might use this information to support his belief in the existence of a target, or, perhaps, determine what strategy to pursue.

b. Rule-Based System Directs Signal Processing Devices

The expert system proposed would interface with the signal processing devices on both the data and control level. The rule-based system could, like the human expert, direct the behavior of these devices and use their output to trigger other knowledge sources. This could save the domain expert and the rule-based system time while making intelligent use of more efficient devices.

6. Limitations of the Artificial Intelligence System

The primary goal of the system is to come to as complete an understanding of the principals and factors which affect detecting, tracking, and identifying the signals in the presence of noise. The goal of the rule-based system is to aid in that endeavor. It is good to be aware of the capabilities and limitations of artificial intelligence machines.

a. Symbolic Manipulation Computers Are Not Comparatively Fast

One of the limitations of symbolic computers is that they are not particularly fast when compared to most signal processing machines. So if real time operation is an eventual goal, the successful artificial intelligence routines could be replaced with algorithms that run on faster general or special purpose computers. Parallel processing should be considered if alternative solutions are possible. Compiling the artificial intelligence routines is another option that can be used to speed execution.

b. Lacks Breadth and Depth of Knowledge

Depth and breadth of knowledge play an important part in the ability of the domain expert. It is difficult to transfer this to a rule-based system, thus, the rule-based system responds more like an *idiot savant*, remarkably capable in an area but unable to recognize when its knowledge is insufficient or inadequate.

c. Lacks the Ability to Independently Check Results

Expert systems have no independent way of checking their solution. For example, the computer has no way of positively knowing where an animal is. If the computer reports it is in a specific area,

periodic verification checks by independent means will aid in system development and subsequently build confidence in the system. False alerts can be reduced if this type of feedback is provided.

C. ARTIFICIAL INTELLIGENCE TECHNIQUES

Four artificial intelligence areas were determined to have the greatest potential benefit for use in a rule-based signal interpretation system. They are: reasoning with uncertainty; the frames data structure, which supports pattern matching; the blackboard architecture that can support both uncertainty, and the frames structure; and learning systems.

1. System Uncertainty

Three general sources of uncertainty, identified by Cohen, (Cohen, 1984, p. 29), are the data or evidence, the model, and the beliefs or results. Uncertainty in beliefs and results follows from dealing with the uncertainty in the data and model.

a. Data Uncertainty

The type of sensory input being dealt with has two characteristics which make conclusions about it difficult: the target signal is not very large in amplitude, and generally the noise mixed with it is higher in amplitude. This noise can, and most often does, interfere with the target signals. Thus, the expert, and the artificial intelligence sub-system must deal with data uncertainty.

b. Model Uncertainty

The problem being dealt with has no simple algorithmic solution. Expert problem solving strategy generally involving heuristics that give acceptable results most of the time.

Using a heuristic implies uncertainty. As indicated by Cohen and Gruber, (Cohen, 1984, pp. 27-29), expert inference rules are compilations of dozens or even hundreds of experiences and that minor differences are *smoothed out* in the rule. This *smoothing out* means that the rule is statistically correct, i.e., more often correct than not, and is subjectively based on the judgment of the expert.

c. Methods to Deal With Uncertainty

Three approaches to uncertainty are discussed by Cohen and Gruber, (Cohen, 1984, pp. 29-33): the engineering approach, which tries to circumvent uncertainty; the control approach, that tries to avoid uncertainty by selecting paths that lead to more certainty; and a hybrid of the two approaches.

The engineering approach is rejected for knowledge sources that work with noisy data from sensors, and heuristics. This leaves the control approach which suggests use of an artificial intelligence architecture which emphasizes control. The blackboard architecture has strong emphasis on control and is discussed later in this section.

Methods of dealing with uncertainty range from the Baysian approach suggested by Cheeseman, (Cheeseman, 1984, pp. 115-121), to ones that use explicit statements about the uncertainty of their conclusions by Weiss and Kulikowski, (Weiss, 1984, p. 26). Most of the methods cite the need to keep track of where the uncertainty comes from. Debugging symbolic reasoning programs that use uncertainty is considerably harder because of the difficulty in seeing how all possible combinations of rules are affected by the numerical values assigned to production rules.

There have been many papers written dealing with uncertainty in artificial intelligence systems. One source that presents several approaches is Rowe, (Rowe, 1986, Chapt. 8). A few of the methods used are boolean evaluation, confidence factors (fuzzy logic), and nearest neighbor (closest match).

2. The Blackboard Architecture

a. Motivation for the Blackboard Architecture

It is highly desirable to chose an artificial intelligence architecture that supports and facilitates the use of multiple knowledge sources. It is important because the system will surely expand and evolve as new knowledge is sought. The architecture should also provide conflict resolution between competing solutions, and the mechanism for the knowledge sources to communicate, share results, and remain relatively independent. The architecture must be responsive to the new data, i.e., *data* or *event* driven. The blackboard architecture provides these capabilities.

b. Blackboard Architecture Is Similar to Brainstorming

The blackboard architecture works much like a group of experts in a brainstorming session. The goal of brainstorming is to make progress toward solving a problem.

The group has a leader, who knows the expert's capabilities, hands out tasks to the experts, and checks on their progress. The task leader posts partial results and problem status on the blackboard.

The human experts are knowledge sources in the blackboard architecture. The experts work on their problem, as directed by the task leader, and report back with their results. Experts may report that the

task is completed, or that more information is needed. Each expert works independently but is free to use the results posted on the blackboard. The expert is not confined to a particular mold of problem solving, and is free to chose any available tools.

c. Parts of a Blackboard

The basic parts of the blackboard architecture are knowledge sources, a blackboard for posting the status of the problem, and activation records which determine which knowledge sources are activated. Some knowledge sources may be used for control but the majority are for data manipulation.

(1) <u>The Control Knowledge Source</u>. The blackboard's control system is generally located in a knowledge source but other configurations have been suggested by Hayes-Roth, (Hayes-Roth, 1985, pp. 251-321). The task of the control knowledge source is to map out strategies and influence which knowledge sources get activated. This top level module is aware of such things as: time constraints; the success ratio of a strategy or on a finer level, a module; the total progress made by the knowledge sources towards a solution; what to do if complete solutions are not possible; what assets are available; knowledge source execution times; etc.. Activation of a knowledge source is usually based on the occurrence of new data.

d. Granularity of a Knowledge Source

A blackboard architecture which allows the knowledge source to run to completion is called course grained. With the difficulty of the task being addressed it would be better to select a fine grained blackboard architecture which has the ability to reason with partial results.

(1) Granularity is Problem Dependent. The granularity of the level is entirely problem dependent. The knowledge sources should be somewhat independent but this does not limit their size. It is reasonable to break a function into modules over which control is necessary. For instance, parts of the problem which have numerical solutions, or can be accomplished with signal processing devices, should be modularized. These modules, as discussed earlier, may be done by peripheral devices interfaced to the symbolics computer.

(2) Finer Grain Increases Flexibility. Reasoning with finer grained results makes the system more flexible. Every time a partial result is completed the blackboard regains control and can decide if there is another strategy which would be more profitable to pursue. Partial and complete results are written on the blackboard.

e. Flexibility of the Blackboard Architecture

The flexibility of the architecture is realized at the knowledge sources. The knowledge sources can be written in any language, so long as the interface to the blackboard is maintained. However, if the knowledge source is under development then it is best to use a structure that has a natural language interface.

3. The Frames Data Structure

a. Motivation for Frames Data Structure

There are several ways to describe what using the frames structure does. The frames structure is organized to support the methods which experts are believed to reason, it is, therefore, easy for the domain expert to use. Frames are strongly oriented toward pattern matching, a

technique that is heavily used in passive surveillance problems. Frames support taxonomies, that is classes, and sub-classes.

The frames data structure was suggested by Minsky, (Minsky, 1981, pp. 95 - 128). It offers the experts an easy and powerful way of dealing in plain language, with a symbolic reasoning computer.

b. Frames Structure

The frames data structure, as indicated by Fikes and Kehler, (Fikes, 1985, p. 904), provides the knowledge expert an easy method to express domain knowledge. It provides a structured representation of an object or a class of objects. There are constructs available in a frame language for organizing frames that represent classes into taxonomies.

(1) Frame Nodes. The nodes of the tree serve several important purposes: they provide a way for the expert to store comments in a natural language; the meaning of being at that node; what to do with certain parameters; where to go from the present node; and default values. The nodes also have the ability to store routines for manipulating or describing data at that node called *own slots*, which stores information about itself, and *member slots*, which are used when sub-classes are needed.

(2) Frame Slots. The frames structure provides slots, at the nodes for the storage of facts, production rules, inferences, explanations, etc.. Several kinds of links then provide the mechanism for inheritance between nodes, which completes the structure needed to support taxonomies.

(3) Frame Links. Links are the transition from node to node, member links are used for class membership and sub-class links for

sub-class containment. The later is used in representing taxonomies. Consideration must also be given to the transitions between the levels. These links should be efficient and bidirectional.

(4) Frame Attributes. The *kinds of slots* and *kinds of links* are a form of variable typing, similar to attributes in a relational data base structure. In fact, there is a great deal of similarity between frames and a relational data base. The attributes support and maintain semantic integrity between the frames. Thus, lower levels of abstraction, discussed in Sub-paragraph (1), "Frame Nodes", of this section, can use default values from higher levels, when the need arises.

(5) Frames Store Default Values. The frame structure is vertical in nature but can be horizontal as well. Default values for the object in the nodes, for example, can be stored in parallel structures. In the example of the jungle, the parallel structure might have an instance of a jungle listed.

An instance of a jungle would list everything that a jungle might have: lions, trees, Tarzan, etc..

(6) Frames Useful as Pattern Matching Structure. Historically, knowledge in pattern analysis is partitioned into sub-problems which are represented by taxonomic levels of abstraction. Taxonomic in the sense that instances of the pattern appear on the lowest level and the most abstract description of the pattern at the highest.

An example of this type of structure would be breaking the jungle problem down into sub-classes. The level below jungle would be all the animals in the jungle which made sounds; feline, etc.. The level of

the feline might contain all the different kinds of cats in the jungle. The lowest level would be instances of the actual animal.

4. Learning Knowledge

a. Motivation for Automatic Acquisition of Knowledge

The automatic acquisition of knowledge is highly desirable goal. Loosely interpreted, automatic acquisition of knowledge means the rule-based system gains in ability without the expert present. Presently, there are only a few artificial intelligence systems that have successfully automated the task of knowledge acquisition. The systems which have been successful were written, at great expense, Hayes-Roth, Waterman, Lenat, (Hayes-Roth, 1983, p. 155), for very specific applications.

b. Implementation of Learning

The knowledge source for learning can, for example, be specifically targeted to look at execution time versus success of a knowledge source. It could then learn to by-pass the ones which are slow and have a lower probability of success. It could do this by using a rule that selected a knowledge source only if it had a certain execution time to success ratio.

Alternatively, the learner knowledge source could teach the slow but often successful knowledge source to be more efficient. This could be done by the technique of successive refinement used in the following way: given the knowledge source was successful but slow, the teacher could, in some programmatic manner, successively by-pass one or more of the rules in the module, on a tentative basis, and measure their effect by the execution time verses success ratio.

This successive refinement technique of paring down a knowledge source could conversely be used to add rules to knowledge sources, however, this involves storing all or many of the variables deemed *interesting* by the expert. The *interesting* items could be examined by the method suggested in Paragraph 3 of this Section, "Learning by Data", and rules added as their correlation to the targeted signal source rose above some arbitrary level chosen by the expert.

c. Types of Knowledge Which Can Be Learned

(1) Strategic Knowledge. Strategic knowledge involves the use of heuristics. Heuristics, lack the rigors of proof and fall into the category of - rules of thumb, common practice, common sense, which apply to a given situation in a certain set of circumstances.

(2) Factual Knowledge. Factual knowledge involves rules to follow which are backed by rigorous proof. This, as previously suggested, should be handed off to more efficient devices.

(3) Learning Strategy. Learning strategy is more or less teaching the machine what to do in the presence of certain facts or events - given a certain goal. This method is sometimes called teaching, or learning, by example. It is also the way most higher level creatures are thought to learn.

The learner would be a separate knowledge source that scanned historical event files looking for patterns of strategy used by the operator in the presence of certain events or sequences of events.

(4) Incorporating Strategy. Knowledge sources can be developed that are capable of inferring rules from experience. The new rules can

then be either directly incorporated or displace existing rules in the subject knowledge source.

(5) Frames Supports Learning. The frames data structure supports this type of learning system quite well because of its prototype and taxonometric capabilities. General prototype examples could be set up by the domain and knowledge experts. The general prototype would contain default values, procedures, etc. which get used by members of that class which are learned. Procedures, situations, states, and other pertinent information could also be stored in slots of the frame. New strategies stored in the frames could then be attached by links to the general module and tested.

5. Learning From the Data

Another type of learning system that looks directly at data and by-passes the domain expert is suggested by Cheeseman, (Cheeseman, 1984, pp. 115-121), who observes that the main bottleneck in building expert systems is the time necessary for the expert and the knowledge engineer to find and debug a useful set of rules for a given domain. He further believes that it is possible to by-pass the expert and induce the required information directly from the data. He feels this is the only possible approach in domains that are probabilistic, i.e., where there is no known causal or deterministic theory. In his scheme information is extracted in the form of significant joint probabilities that can then be used to compute the conditional probability of any attribute value of interest given other information about the individual concerned.

The article goes on to describe the rational behind his belief and how to implement such a system. A strong point of the Cheeseman's theory, is the use of Bayesian theory.

6. Obstacles to Learning

A major obstacle in learning systems is their inability to deal with unexpected errors which are sent down to it during a learning session. Another problem of automatic learning is generalization. It is difficult to infer a class of patterns which may have an infinite number of patterns from a finite number of observed patterns.

IV. RECOMMENDATIONS

The selection of artificial intelligence structures is more important then the decision to use artificial intelligence. The knowledge engineer, after thoroughly studying the application should have a good idea of what, generally, is needed.

This study recommends several generic artificial intelligence structures. In each structure, the blackboard, and the frames, there are a few shells available commercially. The exact model or product chosen, unfortunately, depends on whether it is available on the artificial intelligence computer chosen. It might therefore seem logical to choose the exact structure, and then look for the machine. However, other factors not addressed in this paper, such as speed, cost and availability, which affect the overall system, may preclude this approach and dictate a tradeoff.

The symbolic representation language should be carefully selected for understandability, conciseness, efficiency, and a vocabulary rich enough to express everything desired by the experts. It should be equipped with a good editor, debugger, and trace facilities. These tools will expedite development.

The use of the blackboard architecture fulfills several needs of the system. It reasons with partial solutions and multiple knowledge sources. It is an opportunistic architecture, the events which happen during the cycle activate knowledge sources in accordance with the strategy planned by the control blackboard. The blackboard architecture supports incremental development, gives the user feedback which is more

understandable, supports the pursuit of multiple solutions and therefore the use of parallel processing.

The use of knowledge sources has several benefits. Development is easier when the experts knowledge is modularized. It is easier to understand, change, and debug. At run time knowledge sources can be moved in an out of working memory without interference with other modules. Knowledge sources can, where the strategy permits, be run in parallel.

The frames data structure, with its tree like structure, supports classes, sub-classes, and taxonomies and is well suited for pattern matching. Prototypes of both signal and noise sources are easy to access in the frames structure. The frames structure also provides the facility to store, among other things, plain language text, and procedures at the nodes.

Uncertainty occurs in two areas of the system; the data, and the strategy. Several methods are suggested for dealing with these areas.

It is recommended that the rule-based system be interfaced with other computers if real time operation is desired. Thus, the rule-based system is used as the investigative tool that it is, and when suitable algorithms are found, they can be removed them from the rule-based system, and run on more efficient devices. These peripheral devices are controlled by the rule-based system. The rule-based system will run faster, perhaps development will go faster, and the goal of working in real time will be closer.

If the decision is made to use artificial intelligence for the problem, the expectation level of system performance should be set at a realistic

level. If the problem has defied solution for any length of time, it is not reasonable to believe that addition of artificial intelligence to the project will make those go away. The knowledge sources used by the rule-based system may take a long time to develop to the satisfaction of the domain expert.

The rule-based system is being used as a tool to investigate the problem and to gain insight. The tasks it does, it does well, and consistently. The rule-based system can sometimes make inferences that the human was not aware of, because it can put together long chains of inferences. Unlike the human expert, the rule-based system does not know when it does not have enough facts. It cannot tell when it is being asked to reason about something in which it does not have enough knowledge. It is reasonable to assume that these problems can, with time and effort, be overcome.

It is essential to check the results reached by the rule-based system with an independent source. This will be of statistical value in the learning knowledge source and either give the domain expert confidence in the rule-based system or a clue into where to start debugging.

LIST OF REFERENCES

Bolc, L., "Natural Language Interface to the Question-Answering System for Physicians", Computers and Artificial Intelligence, No. 1, March 1984.

Cheeseman, P., "Learning of Expert Systems from Data", Proceedings on Knowledge Based Systems, IEEE Computer Society Press, 1984.

Cohen, P. R., and Gruber, T. R., Expert Systems, Pergamon Infotech Limited, 1984.

Erman, L. D., London, P. E., and Fickas, S. F., "The Design and an Example use of Hearsay-III, Proceedings of the Seventh International Joint Conference on Artificial Intelligence, 24-28 August 1981.

Fikes, R., and Kehler, T, "The Role of Frame-Based Representation in Reasoning", Communications of the ACM, Vol. 28, Number 29, September 1985.

Hayes-Roth, B., "A Blackboard Architecture for Control", Artificial Intelligence an International Journal, Vol. 26, 1985.

Hayes-Roth, F., Waterman, D. A., and Lenat, D. B., Building Expert Systems, Addison-Wesley Publishing Co., 1983.

Kopec, G. E., Oppenheim, A. V., and Davis, R., "Knowledge-Based Signal Processing", Trends and Perspectives in Signal Processing, Vol.2, No.3, July 1982.

Minsky, M., "A Framework for Representing Knowledge", Mind Design, Chapt. 3, MIT Press, 1981.

Rowe, N. C., Introduction to Artificial Intelligence Through Prolog, Prentice Hall, 1987.

Rychener, M. D., Banares-Alcantara, R., and Subrahmanian, E., "A Rule-Based Blackboard Kernel System: Some Principals in Design", Proceedings on Knowledge-Based Systems, IEEE Computer Society Press, 1984.

Weiss, S. M., and Kulikowski, C. A., A Practical Guide to Designing Expert Systems, Rowman and Allanheld, 1984.

INITIAL DISTRIBUTION LIST

		No. Copies
1.	Defense Technical Information Center Cameron Station Alexandria, VA 22304-6145	2
2.	Commander, Space and Naval Warfare Systems Command Mr. Ronald Mitnick PMW 180-421 Washington, DC 20363-5100	1
3.	Commander, Space and Naval Warfare Systems Command PMW 180T Washington, DC 20363-5100	1
4.	Commander, Space and Naval Warfare Systems Command PMW 180-3 Washington, DC 20363-5100	1
5.	Commander, Space and Naval Warfare Systems Command PMW 180-4 Washington, DC 20363-5000	1
6.	Commander, Space and Naval Warfare Systems Command PMW 180-42 Washington, DC 20363-500	1
7.	Commander, Naval Research Laboratory Code 5100 4555 Overlook Ave SW Washington, DC 20375-5000	1
8.	Commander, Naval Research Laboratory Code 5109 4555 Overlook Ave SW Washington, DC 20375-5000	1
9.	Commander, Naval Underseas Systems Center ATTN: Ronald Keipher, Code 3211 New London, CT 06320	1

10. Commander, Naval Air Development Center 1
ATTN: Ed Gerabed, Code 3032
Warminster, PA 01874-5000

11. Commander, Naval Ocean Systems Center 1
ATTN: Tracy Ball, Code 71
San Diego CA 92152-5000

12. Commander, Naval Ocean Systems Center 1
ATTN: John Roese, Code 73
San Diego CA 92152-5000

13. Defense Advanced Research Project Agency 2
1400 Wilson Blvd.
Arlington, VA 22209-2308

14. Commander, Naval Ocean Systems Center 1
ATTN: Mr. Bill Powell, Code 50
San Diego, CA 92152-5000

15. Officer In Charge, Naval Ocean Systems Center 1
ATTN: Mr. Dan Hightower, Code 015
Hawaii Laboratory
PO Box 997
Kailua, HI 96734

16. Commander, Naval Ocean Systems Center 1
ATTN: CDR Steven Ross
San Diego, CA 92152-5000

17. Commander, Naval Ocean Systems Center 1
ATTN: Mr. Ronald Seiple
San Diego, CA 92152-5000

18. Commander, Naval Ocean Systems Center 1
ATTN: Mr. Alfred Nejame, Code 142
San Diego, CA 92152-5000

19. Library, Code 0142 2
Naval Postgraduate School
Monterey, CA 93943-5002

20. Department Chairman, Code 52 1
Department of Computer Science
Naval Postgraduate School
Monterey, CA 93943-5000

21. Curriculum Officer, Code 37 1
Computer Technology Programs
Naval Postgraduate School
Monterey, CA 93943-5000

22. Professor Neil C. Rowe, Code 52 1
Computer Science Department
Naval Postgraduate School
Monterey, CA 93943-5000

23. Professor Calvin R. Dunlap 1
Oceanography Department, Code 68 DU
Naval Postgraduate School
Monterey, CA 93943-5000

24. Officer In Charge, Naval Ocean Systems Center 1
ATTN: Dave Smith, Code 531
PO Box 997
Kailua, HI 96734

25. Officer In Charge, Naval Ocean Systems Center 3
ATTN: Kirk E. Jennings, Code 531
PO Box 997
Kailua, HI 96734

220283

Thesis
J47254
c.1

Jennings
Design guidelines for a rule-based passive surveillance system.